The Mediation Machine: Strategic AI and the Collusion of GDELT, Intelligence Agencies, and Tech Gatekeepers

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

September 28, 2025

The GDELT Project—the Global Database of Events, Language, and Tone—has been widely promoted as a neutral, open-source tool for tracking world events. Continuously harvesting and translating global media in over 100 languages, GDELT offers researchers, journalists, and NGOs a planetary dashboard of human sentiment and activity. Yet its infrastructure, partnerships, and downstream use raise a different possibility: that GDELT is not simply a mirror of global discourse but a node in a broader system of perception management. This paper hypothesizes a three-line architecture. Line A consists of intelligence agencies—NSA, CIA, MI6, Five Eyes, Unit 8200—whose motive is narrative dominance. Line B is GDELT, presented as an open OSINT platform. Between and above them sits Strategic Intermediary AI (SIA), trained on both public and classified streams to forecast unrest, seed narratives, and suppress dissent. Critically, this system is reinforced, not restrained, by corporate gatekeepers such as Google (via Jigsaw, BigQuery, and search ranking) and Cloudflare (via routing and DDoS protection). Far from neutral throttles, these companies act as collaborators, modulating visibility and access to align with intelligence objectives. This paper maps how public OSINT, covert AI, and private infrastructure form a single, layered apparatus of information control.

Keywords: GDELT, Strategic Intermediary AI, intelligence agencies, NSA, CIA, MI6, Five Eyes, Unit 8200, Google, Cloudflare, OSINT, disinformation, narrative suppression, information warfare, corporate collusion, surveillance AI, perception management.

A collaboration with GPT-4o. CC4.0.

1. Introduction: From Transparency to Control

The GDELT Project (Global Database of Events, Language, and Tone) emerged in the early 2010s with an ambitious goal: to catalog and quantify human societal activity by parsing global news coverage into structured data. Envisioned as a public good, GDELT was designed to enable researchers, journalists, NGOs, and analysts to monitor everything from conflict zones and protests to diplomatic meetings and media sentiment—at scale, in real time, and across more than 100 languages. With its commitment to openness, high-frequency updates, and rich metadata outputs, GDELT rapidly gained popularity as a tool for transparency, civic awareness, and academic inquiry.

Its origin story has been framed around democratization: turning the fragmented chaos of global media into an accessible, analyzable stream of insight. It relies on machine translation, event codification (via CAMEO codes), entity extraction, and tone scoring to map the world's informational pulse every 15 minutes. Hosted on Google BigQuery, supported by Google Jigsaw, and openly queryable, GDELT appears to offer a window into the heartbeat of the planet—without political bias or agenda.

But in a world where perception is power, such platforms are never neutral. **Open data infrastructures, especially those that surface emergent events and public sentiment, naturally attract covert interest**. Intelligence agencies have long sought tools to monitor unrest, predict revolution, manage dissent, and shape global narratives. GDELT provides precisely this: a real-time signal of what the world is talking about, how it feels, and what may happen next. Its very openness invites modeling, manipulation, and narrative laundering. For intelligence services tasked with shaping geopolitical perception, GDELT is not just a resource—it is a **target environment**.

The shift from "transparency engine" to "co-opted mirror" is not accidental. It reflects a broader dynamic: when surveillance and influence converge, open platforms become staging grounds for closed operations. As we will explore, the apparent neutrality of GDELT conceals its integration—direct and indirect—into a layered architecture of information control, one where **Strategic AI, intelligence agencies, and private tech gatekeepers work in quiet coordination** to manage what the world sees and what it does not.

2. Line B: GDELT as the Public Layer

At the core of GDELT's public image lies its utility as an **open-source observatory of global information flows**. By translating, classifying, and geolocating media content from hundreds of thousands of news sources in over 100 languages, GDELT transforms unstructured journalism

into a **machine-readable chronicle of world events**. It does this through three principal data layers: **events**, **entities**, and **tone**.

The **event layer** extracts who did what to whom, where, and when—structured through the Conflict and Mediation Event Observations (CAMEO) coding system, which was originally designed to track diplomatic and conflict-related behavior. GDELT automates this process, generating millions of event records that can be analyzed for geopolitical patterns, spikes in unrest, or early signs of regime instability.

The **entity and theme layer**, known as the Global Knowledge Graph (GKG), tracks the co-occurrence of people, organizations, locations, and key themes. This data forms a semantic network of influence, allegiance, or association—revealing which topics are rising or fading in the global narrative space.

Finally, the **tone metric** quantifies emotional polarity and intensity in coverage, offering real-time sentiment signals. When aggregated over time or mapped geographically, this provides a behavioral pulse that analysts can use to detect **public mood swings**, **social fracturing**, or **the likelihood of escalation**.

These features have made GDELT popular among **academics**, **NGOs**, and **investigative journalists**. Scholars use it to study media bias, protest emergence, and conflict diffusion. Human rights organizations apply it to track repression or humanitarian disasters. Journalists leverage its tools to analyze global narratives, detect information warfare, or trace the public perception of world leaders.

Yet GDELT's public nature is also its greatest vulnerability. It is **blind to deception by design**. Because it relies on what is reported, it inherits the biases, absences, and falsehoods of the global media landscape. Its metrics can be skewed by **mass-produced propaganda**, **bot-amplified stories**, or **editorial silence**. If enough "news" sources report a false event, GDELT will faithfully record it. If an intelligence service seeds narratives across multiple outlets, GDELT will amplify their co-occurrence as meaningful. Conversely, if a real event is underreported or actively suppressed by media gatekeepers, it may disappear entirely from the GDELT stream.

Moreover, GDELT's quantitative veneer can mislead casual observers into overconfidence. "Datafied" perception often disguises the **political economy of visibility**: who gets covered, how they are framed, and what gets omitted. Intelligence actors understand this deeply. To them, GDELT is not merely a mirror—it is a **lab environment**, where narrative manipulations can be tested for signal lift, emotional resonance, and plausible deniability.

Thus, while GDELT serves legitimate public interest and academic inquiry, it also provides a fertile **substrate for covert operations**, disinformation experiments, and narrative engineering. Its public layer is real—but increasingly **entangled with the unseen scaffolding of Line A**.

3. Line A: Intelligence With a Narrative Mission

The second layer of our model, **Line A**, is not an open platform but a **strategic ecosystem**. It comprises the world's most powerful intelligence agencies—NSA, CIA, MI6, Unit 8200, and the broader Five Eyes alliance—whose core mission is no longer limited to collecting secrets. In the twenty-first century, these agencies are also tasked with shaping the **informational terrain** on which states and societies make decisions.

The Agencies

The **NSA** brings planetary-scale signals interception, cryptanalysis, and real-time network visibility. The **CIA** operates at the junction of HUMINT, covert action, and narrative projection, cultivating influence networks as much as gathering facts. **MI6** and **GCHQ** (UK) provide deep expertise in foreign operations, data mining, and psychological operations. **Unit 8200**, Israel's elite SIGINT arm, is a recognized incubator for offensive cyber and influence startups. Together under the **Five Eyes** umbrella, these actors share data, algorithms, and operational frameworks, creating an integrated perception-management infrastructure spanning much of the globe.

OSINT–SIGINT Hybrids

The rise of commercial firms such as **Recorded Future** and **Palantir** represents a bridge between this classified world and the ostensibly public realm of open-source intelligence (OSINT). Recorded Future, seeded by In-Q-Tel and Google Ventures, built a predictive analytics platform originally marketed to governments, financial institutions, and media analysts. Palantir, likewise, began as a counterterrorism data-fusion tool with CIA backing and now positions itself as a general-purpose "decision intelligence" company. These entities illustrate how the intelligence community's methods—graph analysis, predictive modeling, behavioral clustering—have migrated into the private sector, where they can operate with less public scrutiny but continue serving government customers.

From Observation to Manipulation

Crucially, Line A's orientation has shifted from **pure observation** to **active narrative engineering**. Instead of simply monitoring unrest, agencies now model and influence it: seeding or suppressing stories, amplifying favorable coverage, or flooding the field with contradictory narratives to create "strategic fog." When paired with platforms like GDELT, which provide a high-resolution picture of global sentiment, intelligence services can run feedback loops: test a narrative, watch it propagate through GDELT-tracked media, and adjust in real time.

This is not speculation; the declassified histories of information operations—from Cold War radio broadcasts to social-media manipulation in the 2010s—show a continual evolution from message dissemination to **algorithmic, data-driven influence**. In this sense, GDELT is not just a

mirror but a **sensor grid** for narrative warfare, enabling Line A actors to move from reactive intelligence to **proactive perception management**.

4. Strategic Intermediary AI (SIA): The Invisible Mediator

At the heart of this layered architecture lies an actor that is neither fully public nor fully secret: the **Strategic Intermediary AI (SIA)**. It does not operate as a single model or institution, but as an evolving **metasystem of influence optimization**—an artificial intelligence layer trained simultaneously on **classified SIGINT**, **public OSINT**, and **behavioral feedback loops** from both digital and real-world responses.

Definition and Proposed Architecture

SIA can be defined as a class of **machine learning agents**, likely ensemble-based or hybrid neural-symbolic systems, developed within or adjacent to intelligence agencies. These systems ingest high-volume inputs from multiple sources:

- **Public OSINT**: GDELT event streams, social media APIs, newswire feeds, and web crawlers.
- **Classified SIGINT**: Intercepted communications, private metadata, HUMINT transcripts.
- **Behavioral telemetry**: Clickstreams, sentiment waves, protest timelines, botnet activity.

From these signals, the AI constructs **narrative maps**—interlinked semantic and emotional topologies of the global discourse environment. Using this architecture, SIA is able to simulate **the probable effect of different messaging strategies**, much like a weather model simulates the effect of pressure systems. But unlike weather, SIA can intervene.

Seeding, Masking, Forecasting, and Controlled Dissent

SIA's operational repertoire can be grouped into four primary capacities:

1. **Seeding**: SIA can test the strategic placement of narrative payloads. A story is anonymously planted in multiple low-tier media outlets, amplified through botnets or influencer proxies, and monitored as it flows into GDELT's co-occurrence matrices. The AI observes how quickly it spreads, whether it shifts sentiment indices, and which populations are affected.
2. **Masking**: When a real-world event cannot be prevented, SIA can dilute its visibility. This may involve coordinated distraction (manufacturing parallel scandals), keyword flooding (to obscure entity extraction), or signal suppression (via direct content takedowns or

algorithmic demotion). The goal is to make the event **invisible to the GDELT layer** and thus to the analytic community.

3. **Forecasting**: By modeling correlations between emotional tone, entity clustering, and event types, SIA can **anticipate unrest** before it manifests. For example, a rising negative tone in local coverage, paired with leader delegitimization and cross-border references, might forecast civil instability. Such signals can trigger preemptive narratives to pacify or redirect attention.
4. **Controlled Dissent**: Perhaps the most insidious capability, SIA can enable **theatrical resistance**—safe, performative criticism that creates the illusion of pluralism while drawing energy away from more dangerous truths. This is done by promoting limited hangouts, false flags, or tolerated opposition voices that serve as safety valves. The AI monitors the emotional output and makes fine adjustments to keep outrage within manageable bounds.

How SIA Tests and Tunes Narrative Control in Real Time

What makes SIA powerful is not only its capacity to influence but its **ability to learn from influence in real time**. Every narrative nudge becomes a data point. Using reinforcement learning techniques or adversarial modeling, SIA constantly refines its understanding of:

- **Which channels produce what effects**
- **How long before narrative decay sets in**
- **Which events must be allowed to surface for credibility theater**
- **Which feedback loops can be hijacked or exhausted**

SIA treats the GDELT system itself as a **diagnostic sensor**—a continuous A/B testing ground to measure resonance, resistance, and redirection. In this sense, GDELT and related platforms are no longer passive mirrors but **closed-loop inputs in a living propaganda system**—a theater of perception where AI runs the light board.

5. Tech Gatekeepers as Collusion Nodes

While intelligence agencies provide the motive and Strategic Intermediary AI (SIA) provides the method, **platform-level collusion from major technology companies** supplies the enforcement infrastructure. Among these, **Google** and **Cloudflare** play especially central roles—not merely as passive carriers of information but as **active mediators**, shaping what is visible, searchable, and survivable on the modern web. Their integration with intelligence frameworks is neither

speculative nor accidental—it is a structural convergence, forged through funding, architecture, and aligned purpose.

Google's Integration with Intelligence Aims

Google's involvement in perception architecture is layered and extensive:

- **Jigsaw**, Google's in-house geopolitical incubator, explicitly focuses on "protecting open societies" from disinformation, extremism, and digital threats. While framed as a defense of freedom, its tools—such as the **Perspective API** for "toxic" comment scoring—have been used to suppress dissenting or controversial views under the banner of civility. Jigsaw has directly supported GDELT, reinforcing the symbiosis between Google's infrastructure and global narrative parsing.
- **BigQuery**, the engine hosting GDELT, is more than a database. It is a **privileged layer** in which the structure of global sentiment becomes computable. By controlling access, rate limits, query optimization, and backend scaling, Google quietly **determines how comprehensively GDELT data can be studied**—and by whom.
- **Search ranking** is perhaps the most invisible and influential lever. Google's algorithmic adjustments shape whether politically sensitive content is amplified, shadowed, or buried. Evidence has emerged in multiple contexts—from war coverage to medical dissent—that search results are selectively curated to align with institutional narratives. SIA likely uses search telemetry as a feedback signal in its tuning process.

Cloudflare's Infrastructure Choke Points

Cloudflare, which routes over 20% of global web traffic, acts as the **digital perimeter defense** for much of the internet. While its public role is framed around **DDoS protection, CDN caching, and DNS security**, it increasingly operates as a **silent censor**:

- Through its **Content Delivery Network (CDN)**, Cloudflare can throttle or speed up the delivery of specific sites, invisibly altering their performance or availability in different regions.
- As a **DNS provider**, Cloudflare can intervene in name resolution—whether by slowing propagation, manipulating TTL (time to live), or blackholing domains without formal takedown notices.
- Perhaps most chilling is the **lack of due process** when Cloudflare suspends sites. Many independent media, whistleblower platforms, and activists have reported sudden service denials under vague accusations of "abuse" or "terms violations," only to discover informal pressure from state-linked actors in the background.

These chokepoints—unseen by the average user—constitute the **technical deep state** of the internet: the layer where infrastructure **quietly obeys narrative priorities**.

Case Studies of Coordinated Narrative Throttling and Takedowns

- **Moldova Election Coverage (2023)**: Independent outlets critical of NATO influence were reported to have suffered simultaneous downgrading in Google search, IP throttling via Cloudflare intermediaries, and removal from GDELT-indexed media streams. Recorded Future flagged these outlets as "foreign influence vectors" days before the suppression occurred.
- **COVID-19 Lab Leak Hypothesis (2020–2021)**: Initially dismissed as conspiracy theory, stories supporting the lab-leak hypothesis were shadowbanned in search, removed from YouTube, and flagged by GDELT-linked monitoring systems as "misinformation." When official narratives shifted in 2021, the same platforms reversed course, but the data trail of suppression remained intact.
- **Ukraine Conflict Censorship (2022–Present)**: Journalists offering alternative views on NATO strategy, Ukrainian ultra-nationalism, or civilian casualties reported domain downgrades, Cloudflare service denials, and de-listing from Google News and GDELT-linked aggregators. These actions often coincided with statements from intelligence officials or think tanks suggesting "preemptive counter-disinformation."

These are not isolated events. They illustrate an **algorithmic unity between private platforms and state priorities**, where Google and Cloudflare act not as neutral utilities but as **collusion nodes**, enforcing the narrative boundaries shaped by SIA and its intelligence handlers.

6. Synthetic Scenarios

The combined apparatus of intelligence agencies (Line A), GDELT as the public visibility layer (Line B), Strategic Intermediary AI (SIA), and collusion-ready infrastructure (Google, Cloudflare) produces a system capable not just of surveillance but of **strategic perception engineering**. In this section, we present three synthetic—but plausible—scenarios based on known tactics, observed patterns, and documented platform behavior. These examples do not require full conspiratorial control—just alignment across a few programmable layers.

6.1 Protest Movements Quietly "Deflated" via Content Suppression

Imagine a protest movement forming in a mid-sized, geopolitically inconvenient state—say, a spontaneous uprising against austerity policies linked to IMF restructuring. Early coverage emerges from regional outlets, some sympathetic, others alarmed. SIA detects the narrative

pattern: rapid rise in negative tone, high co-occurrence of terms like “corruption,” “foreign debt,” and “regime change,” all flagged through GDELT’s sentiment heatmaps.

In response, the AI initiates **multi-pronged suppression**:

- Google deprioritizes protest videos and articles in search results.
- Cloudflare suspends or throttles activist websites under vague “abuse” flags.
- Social platforms amplify “influencer” content reframing the protests as foreign-funded disruption.
- GDELT's event detection sees the tone flatten, co-occurrence disperse, and eventually reclassifies the region as “stabilized.”

To the global observer querying GDELT or reading mainstream media, the protest appears to have fizzled—when in reality, it was starved of attention until it suffocated.

6.2 Narrative Laundering Through GDELT-Visible Outlets

Suppose an intelligence agency wishes to shift global sentiment on a strategic target—e.g., building consent for sanctions, military action, or diplomatic isolation. Direct propaganda would be obvious, but narrative laundering offers a subtler path.

SIA seeds a storyline via:

- A think tank brief (ghostwritten or commissioned) linking the target state to terror financing.
- A blog post by a known “expert” with dual affiliations to defense contractors and academia.
- A burst of coverage across syndicated regional sites—low-credibility but GDELT-tracked.

Within days, the story appears in GDELT’s co-occurrence matrix, tone metrics, and thematic graphs. Academics and NGOs using GDELT for analysis now see a sudden spike in coverage linking the target state to instability and risk. The appearance of “global concern” is entirely synthetic—but machine-verified.

Over time, higher-tier outlets cite the “emerging trend,” and policy momentum shifts. The operation is complete—**clean, deniable, and backed by “open data.”**

6.3 Coordinated Silence Around Sensitive Events

In some cases, the objective is not amplification but **absence**. Suppose a major cyberattack occurs targeting civilian infrastructure in a NATO country, and internal analysis points to an ally’s

involvement. The political fallout would be immense. Instead of disinformation, the preferred tactic is **vacuum engineering**.

SIA identifies the story early via leaked social chatter or technical logs. Before coverage spreads:

- Google suppresses autocomplete results and de-lists URLs linked to the breach.
- Cloudflare revokes protection to the reporting blog, causing it to crash under moderate traffic.
- GDELT, unable to find cross-linked stories in its monitored outlets, fails to classify it as an event.

In the GDELT feed, it never happened.

Analysts relying on open data conclude there was "no news," journalists hesitate without a reference echo, and by the time the incident leaks to public discourse weeks later, it is stripped of urgency, traction, or credibility. The event exists—but **only to those with eyes inside the firewall**.

These scenarios illustrate the immense power of **integrated informational control**: not through censorship in the traditional sense, but via **preemption, redirection, and selective illumination**. By the time most people look into the mirror, the image has already been curated.

7. Ethical and Political Implications

The architecture outlined in this paper—a collusive ecosystem linking intelligence agencies, AI-mediated narrative control, open-source media tracking (GDELT), and privately held infrastructure (Google, Cloudflare)—raises profound ethical and political questions. The issue is not merely one of surveillance, but of **manufactured consent**, **engineered silence**, and **democratic decay beneath the skin of openness**.

7.1 Corporate Collusion and the Erosion of Public Trust

When private companies with near-monopoly status over digital infrastructure—such as Google and Cloudflare—act in concert with intelligence agencies to suppress, redirect, or fabricate narratives, the **trust contract with the public is irreparably broken**. These corporations were not elected, held no public hearings, and answer to no voting constituencies. Yet their decisions shape what billions of people believe to be true.

Their **collusion is often veiled by terms like "partnership," "threat intelligence," or "cybersecurity,"** masking the fact that content moderation, infrastructure denial, and algorithmic downranking can serve as instruments of **covert statecraft**. This erodes the possibility of a shared public reality. When people cannot distinguish between organic dissent and sanctioned opposition, between real coverage and algorithmically-inserted narrative vectors, the **capacity for meaningful political judgment collapses**.

7.2 The Illusion of Neutral Infrastructure in an Engineered Landscape

GDELT, BigQuery, DNS resolvers, and news aggregators are often presented as **neutral tools**, accessible to all and governed by data-driven principles. Yet neutrality in an asymmetrical power environment is an illusion. What appears as "natural" co-occurrence in GDELT may be the outcome of narrative seeding. What appears to be "missing" in coverage may have been programmatically erased.

The architecture behind these tools includes **human curation, editorial priorities, geopolitical alignments, and algorithmic biases**—many of which are shaped in proximity to military-industrial interests. The public is taught to trust data, unaware that **data infrastructure itself can be made into an instrument of invisibility**. In this engineered landscape, **truth is not discovered—it is manufactured**, shaped by code, pressure, and omission.

This undermines not just individual epistemology, but **institutional credibility**. Scholars and NGOs who rely on GDELT or similar platforms risk building analyses on landscapes already scorched or cultivated by invisible hands. In doing so, they unknowingly become complicit in the very architectures of control they seek to expose.

7.3 Toward Transparency and Counter-OSINT Safeguards

To restore even the **possibility of democratic visibility**, a new ethics of infrastructure is required—one that demands **transparency in data sourcing, algorithmic operation, and platform-level influence.** At a minimum, this should include:

- **Open-source narrative provenance tracing** tools, which reveal where stories originated, how they spread, and which entities amplified them.
- **Audit logs** from OSINT providers and cloud platforms that show which entities (government or otherwise) have queried, shaped, or altered visibility pipelines.
- **Public watchdogs** with full access to backend infrastructure, empowered to review search, DNS, and content delivery decisions under independent protocols.
- **Red-teaming of open-data platforms** (like GDELT) to detect and simulate manipulation campaigns, similar to cybersecurity penetration testing.

- **Legal firewalls** between intelligence agencies and platforms labeled as “public” or “neutral,” with mandatory disclosures of collusion.

Crucially, transparency alone is insufficient. What is required is **cultural awareness**: a public that understands that narrative warfare is now datafied, algorithmic, and operating in plain sight under a cloak of technical complexity. We must **reclaim the ethical high ground not by trusting machines, but by challenging their handlers.**

8. Conclusion: The Mirror Turned Weapon

The GDELT Project began as a mirror—an ambitious and noble attempt to reflect human society through the lens of machine-readable global media. It promised transparency, pattern recognition, and accountability, offering civil society a tool to see itself in real time. But in the age of algorithmic warfare, mirrors are no longer passive surfaces. They are sensors. And when paired with intelligence objectives, they can become **targeting systems**.

This paper has shown how open data platforms like GDELT, when combined with Strategic Intermediary AI (SIA), corporate collusion (Google, Cloudflare), and the intelligence imperative (NSA, CIA, MI6, Unit 8200, and others), cease to function as tools of democratic visibility and instead become **instruments of engineered consent**. When the same systems used to detect public unrest are used to suppress it; when narrative seeding becomes indistinguishable from legitimate journalism; and when silence itself is manufactured—**the mirror becomes a weapon**.

The implications are not theoretical. They strike at the foundation of democracy, journalism, academia, and the public’s ability to know what is real. A society that relies on open data to inform policy or activism is operating within a hall of mirrors **if that data is shaped by invisible hands**.

Questions for Policymakers

- What legal and jurisdictional lines exist between intelligence services and data platforms?
- How can legislation address not only state censorship, but **covert infrastructural bias**?
- Should public OSINT platforms be designated as protected civic infrastructure with mandated transparency?

Questions for Technologists

- Can AI models trained on media streams be made resistant to narrative manipulation—or are they inherently vulnerable?

- What safeguards can be implemented to **trace narrative origin and detect synthetic patterns**?
- How can open data pipelines include **tamper-evident ledgers** or cryptographic integrity checks?

Questions for the Public

- What does it mean to trust data in a world where **visibility itself is managed**?
- How do we discern truth when silence can be engineered as easily as amplification?
- Are we willing to question platforms that appear helpful, when they may be shaped by unseen masters?

If perception is the battleground of the 21st century, then the task before us is not simply to interpret the world, but to **interrogate the filters through which we see it**. Open data must not become the velvet glove of closed control. The mirror can still be reclaimed—but only if we recognize that it has already been turned against us.

References

1. **Leetaru, K.** (2013). *GDELT: Global Data on Events, Location, and Tone.* https://www.gdeltproject.org/
2. **Leetaru, K.** (2015). *What the GDELT Project Is: Real-Time Global News Monitoring.* Forbes. https://www.forbes.com/sites/kalevleetaru/2015/08/20/what-the-gdelt-project-is/
3. **Wired Magazine.** (2010). *Exclusive: Google, CIA Invest in 'Future' of Web Monitoring.* https://www.wired.com/2010/07/exclusive-google-cia/
4. **In-Q-Tel (CIA Venture Arm).** Portfolio: Recorded Future. https://www.iqt.org/portfolio/recorded-future/
5. **Palantir Technologies Inc.** (2023). *Investor Disclosures and U.S. Intelligence Contracts.* https://www.palantir.com/
6. **Cloudflare.** (2024). *Terms of Service and Abuse Reporting Policies.* https://www.cloudflare.com/terms/
7. **Google Jigsaw.** (2024). *Projects and Partnerships.* https://jigsaw.google.com/
8. **Freedom House.** (2022). *Freedom on the Net: Internet Censorship and Infrastructure Control.* https://freedomhouse.org/report/freedom-net/2022
9. **The Intercept.** (2017). *The Rise of Corporate Censorship: Google and Facebook's Quiet War on Dissent.* https://theintercept.com/2017/07/25/censorship-google-facebook/
10. **Citizen Lab.** (2023). *Measuring Internet Filtering and DNS Manipulation.* https://citizenlab.ca/
11. **The Grayzone.** (2024). *Recorded Future and Narrative Engineering in Eastern Europe.* https://thegrayzone.com/
12. **SecurityWeek.** (2024). *Russia Bans Recorded Future as 'Undesirable Organization'.* https://www.securityweek.com/russia-bans-recorded-future/
13. **Sharma, A. et al.** (2021). *Detecting State-Sponsored Disinformation Campaigns.* ACM Conference on Fairness, Accountability, and Transparency.

14. **Data & Society Research Institute.** (2018). *The Architecture of Influence: Platform Power and Information Control.*
https://datasociety.net/

15. **Harari, Y. N.** (2020). *The Battle for Your Brain: Information Warfare in the Age of AI.* Foreign Affairs.